AMENDMENTS TO THE CLAIMS

This listing of claims will replace all prior versions and listings of claims in the application.

1   1.   (Currently amended) A method of selecting a heuristic class for data placement in
2   a distributed storage system comprising the steps of:
3   forming an integer program for each of a plurality of heuristic classes,
4   each of the heuristic classes providing a technique for placing data within the
5   distributed storage system, each of the integer programs comprising an
6   objective of minimizing a replication cost for placing the data;
7   solving each of the integer programs which provide the replication cost for
8   each of the heuristic classes; and
9   selecting the heuristic class having a low replication cost.

1   2.   (Currently amended) A method of selecting a heuristic class for data placement in
2   a distributed storage system comprising the steps of:
3   forming a general integer program which models placing data within the
4   distributed storage system; the data placement;
5   forming a specific integer program which models a heuristic class for the
6   data placement, that provides a technique for placing the data within the
7   distributed storage system, the general and specific integer programs each
8   comprising an objective of minimizing a replication cost for placing the data;
9   solving the general integer program which provides a general lower bound
10   for the replication cost;
11   solving the specific integer program which provides a specific lower
12   bound for the replication cost; and
13   selecting the heuristic class if a difference between the general lower
14   bound and the specific lower bound is within an allowable amount.

1   3.   (Original) The method of claim 2 wherein inputs used in the steps of forming the
2   general and specific integer programs comprise a system configuration, a workload,

3 and a performance requirement.

1 4. (Original) The method of claim 3 wherein the performance requirement comprises

2 a bi-modal performance metric.

1 5. (Original) The method of claim 4 wherein the bi-modal performance metric

2 comprises a criterion and a ratio of successful attempts to total attempts.

1 6. (Original) The method of claim 3 wherein the performance requirement comprises

2 a data access latency.

1 7. (Original) The method of claim 3 wherein the performance requirement comprises

2 a data access bandwidth.

1 8. (Original) The method of claim 3 wherein the performance requirement comprises

2 a data update time.

1 9. (Original) The method of claim 3 wherein the performance requirement comprises

2 an average data access latency.

1 10. (Original) The method of claim 3 wherein the performance requirement comprises

2 a data availability requirement.

1 11. (Original) The method of claim 3 wherein the general integer program comprises

2 general constraints which model the data placement irrespective of the heuristic class

3 for the data placement.

1 12. (Original) The method of claim 11 wherein the general constraints comprise a

2 performance constraint which models the performance requirement.

1 13. (Original) The method of claim 11 wherein the specific integer program

2   comprises the general constraints and a specific constraint.

1   14.   (Original) The method of claim 12 wherein the specific constraint comprises a
2   storage constraint.

1   15.   (Original) The method of claim 12 wherein the specific constraint comprises a
2   replica constraint.

1   16.   (Original) The method of claim 12 wherein the specific constraint comprises a
2   routing knowledge constraint and further wherein the routing knowledge constraint
3   models an extent to which a data storage node knows of replicas of data objects
4   stored on other data storage nodes.

1   17.   (Original) The method of claim 12 wherein the specific constraint comprises an
2   access knowledge constraint and further wherein the access knowledge constraint
3   models an extent to which a data storage knows of access to replicas of data objects
4   by clients accessing other data storage nodes.

1   18.   (Original) The method of claim 12 wherein the specific constraint comprises an
2   activity history constraint.

1   19.   (Original) The method of claim 12 wherein the specific constraint comprises a
2   reactive placement constraint.

1   20.   (Original) The method of claim 3 wherein the system configuration comprises a
2   plurality of data storage nodes coupled by a plurality of network links.

1   21.   (Original) The method of claim 20 wherein the system configuration further
2   comprises a plurality of clients coupled to the data storage nodes.

1   22.   (Original) The method of claim 21 wherein the workload comprises at least some

2   of the clients requesting data objects stored on the data storage nodes.

1   23.   (Original) The method of claim 22 wherein the workload further comprises at

2   least some of the clients storing some of the data objects on the data storage nodes.

1   24.   (Currently amended) A method of selecting a heuristic class for data placement

2   in a distributed storage system comprising the steps of:

3   forming a general integer program which models placing data within the

4   distributed storage system; the data placement;

5   forming a plurality of specific integer programs which model a plurality of

6   heuristic classes, each of the heuristic classes providing a technique for

7   placing the data within the distributed storage system, the general and specific

8   integer programs each comprising an objective of minimizing a replication

9   cost for placing the data;

10   solving the general integer program which provides a lower bound for the

11   replication cost;

12   solving the specific integer programs which provides the replication cost

13   for each of the heuristic classes; and

14   selecting a particular heuristic class correlated to a low replication cost if a

15   difference between the lower bound and the low replication cost is within an

16   allowable amount.

1   25.   (Currently amended) A computer readable memory comprising computer code for

2   implementing a method of selecting a heuristic class for data placement in a

3   distributed storage system, the method of selecting the heuristic class comprising the

4   steps of:

5   forming an integer program for each of a plurality of heuristic classes,

6   each of the heuristic classes providing a technique for placing the data within

7   the distributed storage system, each of the integer programs comprising an

8   objective of minimizing a replication cost for placing the data;

9   solving each of the integer programs which provide the replication cost for

10   each of the heuristic classes; and

11   selecting the heuristic class having a low replication cost.

1   26.   (Currently amended) A computer readable memory comprising computer code for

2   implementing a method of selecting a heuristic class for data placement in a

3   distributed storage system, the method of selecting the heuristic class comprising the

4   steps of:

5   forming a general integer program which models placing data within the

6   distributed storage system; the data placement;

7   forming a specific integer program which models a heuristic class for the

8   data placement, that provides a technique for placing the data within the

9   distributed storage system, the general and specific integer programs each

10   comprising an objective of minimizing a replication cost for placing the data;

11   solving the general integer program which provides a general lower bound

12   for the replication cost;

13   solving the specific integer program which provides a specific lower

14   bound for the replication cost; and

15   selecting the heuristic class if a difference between the general lower

16   bound and the specific lower bound is within an allowable amount.

1   27.   (Currently amended) A computer readable memory comprising computer code for

2   implementing a method of selecting a heuristic class for data placement in a

3   distributed storage system, the method of selecting the heuristic class comprising the

4   steps of:

5   forming a general integer program which models placing the data within

6   the distributed storage system; the data placement;

7   forming a plurality of specific integer programs which model a plurality of

8   heuristic classes, each of the heuristic classes providing a technique for

9   placing the data within the distributed storage system, the general and specific

10   integer programs each comprising an objective of minimizing a replication

11   cost for placing the data;

12   solving the general integer program which provides a lower bound for the

13   replication cost;

14   solving the specific integer programs which provides the replication cost

15   for each of the heuristic classes; and

16   selecting a particular heuristic class correlated to a low replication cost if a

17   difference between the lower bound and the low replication cost is within an

18   allowable amount.